



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/625,320	07/23/2003	David A. Brown	2037.2005-008	7755

21005 7590 06/16/2005

HAMILTON, BROOK, SMITH & REYNOLDS, P.C.
530 VIRGINIA ROAD
P.O. BOX 9133
CONCORD, MA 01742-9133

EXAMINER

MOFIZ, APU M

ART UNIT PAPER NUMBER

2165

DATE MAILED: 06/16/2005

Please find below and/or attached an Office communication concerning this application or proceeding.

RECEIVED
OIPE/IAP

JUN 22 2005

Office Action Summary	Application No. 10/625,320	Applicant(s) BROWN, DAVID A.	
	Examiner Apu M. Mofiz	Art Unit 2165	

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If the period for reply specified above is less than thirty (30) days, a reply within the statutory minimum of thirty (30) days will be considered timely.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

- 1) ☒ Responsive to communication(s) filed on 02 May 2005.
- 2a) ☐ This action is FINAL. 2b) ☒ This action is non-final.
- 3) ☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

- 4) ☒ Claim(s) 1-18 is/are pending in the application.
- 4a) Of the above claim(s) _____ is/are withdrawn from consideration.
- 5) ☐ Claim(s) _____ is/are allowed.
- 6) ☒ Claim(s) 1-18 is/are rejected.
- 7) ☐ Claim(s) _____ is/are objected to.
- 8) ☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

- 9) ☐ The specification is objected to by the Examiner.
- 10) ☒ The drawing(s) filed on 02 May 2005 is/are: a) ☒ accepted or b) ☐ objected to by the Examiner.
Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).
Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).
- 11) ☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

- 12) ☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).
- a) ☐ All b) ☐ Some * c) ☐ None of:
1. ☐ Certified copies of the priority documents have been received.
2. ☐ Certified copies of the priority documents have been received in Application No. _____.
3. ☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).
- * See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

- 1) ☒ Notice of References Cited (PTO-892)
- 2) ☐ Notice of Draftsperson's Patent Drawing Review (PTO-948)
- 3) ☐ Information Disclosure Statement(s) (PTO-1449 or PTO/SB/08)
Paper No(s)/Mail Date _____.
- 4) ☐ Interview Summary (PTO-413)
Paper No(s)/Mail Date. _____.
- 5) ☐ Notice of Informal Patent Application (PTO-152)
- 6) ☐ Other: _____.

DETAILED ACTION

1. Applicant's arguments filed May 02, 2005 with respect to claims 1-18 have been considered but are moot in view of the new ground(s) of rejection.

Claim Rejections - 35 USC § 102

2. The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office action:

A person shall be entitled to a patent unless –

(a) the invention was known or used by others in this country, or patented or described in a printed publication in this or a foreign country, before the invention thereof by the applicant for a patent.

3. Claims 1-5, 8-14 and 17-18 are rejected under 35 U.S.C. 102(a) as being anticipated by Suba Varadarajan (Virtual Local Area Networks, August 14, 1997, pages 1-14 and Varadarajan hereinafter).

As to claims 1 and 10, Varadarajan teaches a switch (page 1) for receiving and forwarding data packets comprising: at least one ingress port for receiving data packets, at least one of the data packets associated with a virtual LAN identifier (page 6); at least one egress port for forwarding the data packets (page 6); and forwarding logic which provides a translated identifier for the virtual LAN identifier (page 9; page 10), the forwarding logic comprising: a filtering database that provides a forward vector for the translated identifier from a single search of the filtering database (page 9; page 10).

Art Unit: 2165

As to claims 2 and 11, Varadarajan teaches wherein the forwarding logic assigns a default value to the virtual LAN identifier (i.e., for all of the VLAN-unaware bridges/switches, the frames are tagged with a default value or untagged. VLMP (i.e., Virtual LAN Management Protocol; for further clarification See VLMP RFC) also requires that.) (page 8; page 9).

As to claims 3 and 12, Varadarajan teaches that the translated identifier includes a group identifier and a group member identifier for the virtual LAN (page 10).

As to claims 4 and 13, Varadarajan teaches the filtering database includes a first entry and a second entry (i.e., the filtering database includes a plurality of entries) (pages 9-11).

As to claims 5 and 14, Varadarajan teaches that the filtering database provides the forward vector associated with the first entry when the translated identifier matches the first entry and provides the forward vector associated with the second entry when only the group identifier portion of the translated identifier matches the second entry (i.e., for each entry, there is a list of entries that the packet can be forwarded to as specified in the filtering database) (pages 9-14).

As to claims 8 and 17, Varadarajan teaches that the virtual LAN identifier is stored, in a header included in the received data packet (pages 9-14).

As to claims 9 and 18, Varadarajan teaches that the forwarding logic further comprises: a forward vector table which stores the forward vector associated with the received data packet, the filtering database providing a pointer to the location of the forward vector in the forward vector table (i.e., the Virtual LAN Management Protocol requires the switch to store a pointer to a list of the Virtual LANs that the source address is in. See VLMP RFC. Explaining RFC is beyond the scope of this Office Action.) (pages 9-14).

Claim Rejections - 35 USC § 103

4. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

5. Claims rejected under 35 U.S.C. 103(a) as being unpatentable over Suba Varadarajan (Virtual Local Area Networks, August 14, 1997, pages 1-14 and Varadarajan hereinafter) in view of Devashish Paul (Ternary DRAM CAM: Now and the Future, pages 1-2 and Paul hereinafter).

As to claims 6-7 and 16-17, Varadarajan teaches a filtering database in a switch (pages 9-14).

Art Unit: 2165

Varadarajan does not explicitly teach that the switch uses Ternary Content Addressable Memory.

It would have been obvious to a person of ordinary skill in the art at the time of Applicant's invention to modify the teachings of Varadarajan with the teachings of Paul to include that the switch uses Ternary Content Addressable Memory with the motivation to use table-lookups at wire speed, conducting as many as 6 million searches per second (Paul, page 1).

Points of Contact

6. Any inquiry concerning this communication or earlier communications from the examiner should be directed to Apu M. Mofiz whose telephone number is (571) 272-4080. The examiner can normally be reached on Monday – Thursday 8:00 A.M. to 4:30 P.M.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Jeffrey Faffin can be reached at (571) 272-4146. The fax numbers for the group is (703) 872-9306.

Any inquiry of a general nature or relating to the status of this application should be directed to the Group receptionist whose telephone number is (703) 305-9600.



Apu M. Mofiz
Patent Examiner
Technology Center 2100

June 13, 2005

Notice of References Cited

Application/Control No.

10/625,320

Applicant(s)/Patent Under

Reexamination

BROWN, DAVID A.

Examiner

Apu M. Mofiz

Art Unit

2165

Page 1 of 1

U.S. PATENT DOCUMENTS

*		Document Number Country Code-Number-Kind Code	Date MM-YYYY	Name	Classification
	A	US-			
	B	US-			
	C	US-			
	D	US-			
	E	US-			
	F	US-			
	G	US-			
	H	US-			
	I	US-			
	J	US-			
	K	US-			
	L	US-			
	M	US-			

FOREIGN PATENT DOCUMENTS

*		Document Number Country Code-Number-Kind Code	Date MM-YYYY	Country	Name	Classification
	N					
	O					
	P					
	Q					
	R					
	S					
	T					

NON-PATENT DOCUMENTS

*		Include as applicable: Author, Title Date, Publisher, Edition or Volume, Pertinent Pages)
	U	Suba Varadarajan, Virtual Local Area Networks, August 14, 1997, pages 1-14
	V	Devashish Paul, Ternary DRAM CAM: Now and Future, 5/1/2000, pages 1-2
	W	
	X	

*A copy of this reference is not being furnished with this Office action. (See MPEP § 707.05(a).)
Dates in MM-YYYY format are publication dates. Classifications may be US or foreign.

Virtual Local Area Networks

Ref 1

Suba Varadarajan, varadarajan.5@osu.edu

This paper describes virtual local area networks (VLAN's), their uses and how they work in accordance with the 802.1Q standard.

[Other Reports on Recent Advances in Networking](#)

[Back to Raj Jain's Home Page](#)

Table of Contents

- [1.0 Introduction](#)
- [2.0 What are VLAN's?](#)
- [3.0 Why use VLAN's?](#)
- [4.0 How VLAN's work](#)
 - [4.1 Types of VLAN's](#)
 - [4.2 Types of Connections](#)
 - [4.3 Frame Processing](#)
- [5.0 Summary](#)
- [6.0 References](#)
- [7.0 Abbreviations](#)

VLAN follows VLMP (Virtual Lan management protocol)
VLMP requires

1.0 Introduction

A Local Area Network (LAN) was originally defined as a network of computers located within the same area. Today, Local Area Networks are defined as a single broadcast domain. This means that if a user broadcasts information on his/her LAN, the broadcast will be received by every other user on the LAN. Broadcasts are prevented from leaving a LAN by using a router. The disadvantage of this method is routers usually take more time to process incoming data compared to a bridge or a switch. More importantly, the formation of broadcast domains depends on the physical connection of the devices in the network. Virtual Local Area Networks (VLAN's) were developed as an alternative solution to using routers to contain broadcast traffic.

In Section 2, we define VLAN's and examine the difference between a LAN and a VLAN. This is followed by a discussion on the advantages VLAN's introduce to a network in Section 3. Finally, we explain how VLAN's work based on the current draft standards in Section 4.

[Back to Table of Contents](#)

2.0 What are VLAN's?

In a traditional LAN, workstations are connected to each other by means of a hub or a repeater. These

devices propagate any incoming data throughout the network. However, if two people attempt to send information at the same time, a collision will occur and all the transmitted data will be lost. Once the collision has occurred, it will continue to be propagated throughout the network by hubs and repeaters. The original information will therefore need to be resent after waiting for the collision to be resolved, thereby incurring a significant wastage of time and resources. To prevent collisions from traveling through all the workstations in the network, a bridge or a switch can be used. These devices will not forward collisions, but will allow broadcasts (to every user in the network) and multicasts (to a pre-specified group of users) to pass through. A router may be used to prevent broadcasts and multicasts from traveling through the network.

The workstations, hubs, and repeaters together form a LAN segment. A LAN segment is also known as a collision domain since collisions remain within the segment. The area within which broadcasts and multicasts are confined is called a broadcast domain or LAN. Thus a LAN can consist of one or more LAN segments. Defining broadcast and collision domains in a LAN depends on how the workstations, hubs, switches, and routers are physically connected together. This means that everyone on a LAN must be located in the same area (see *Figure 1*).

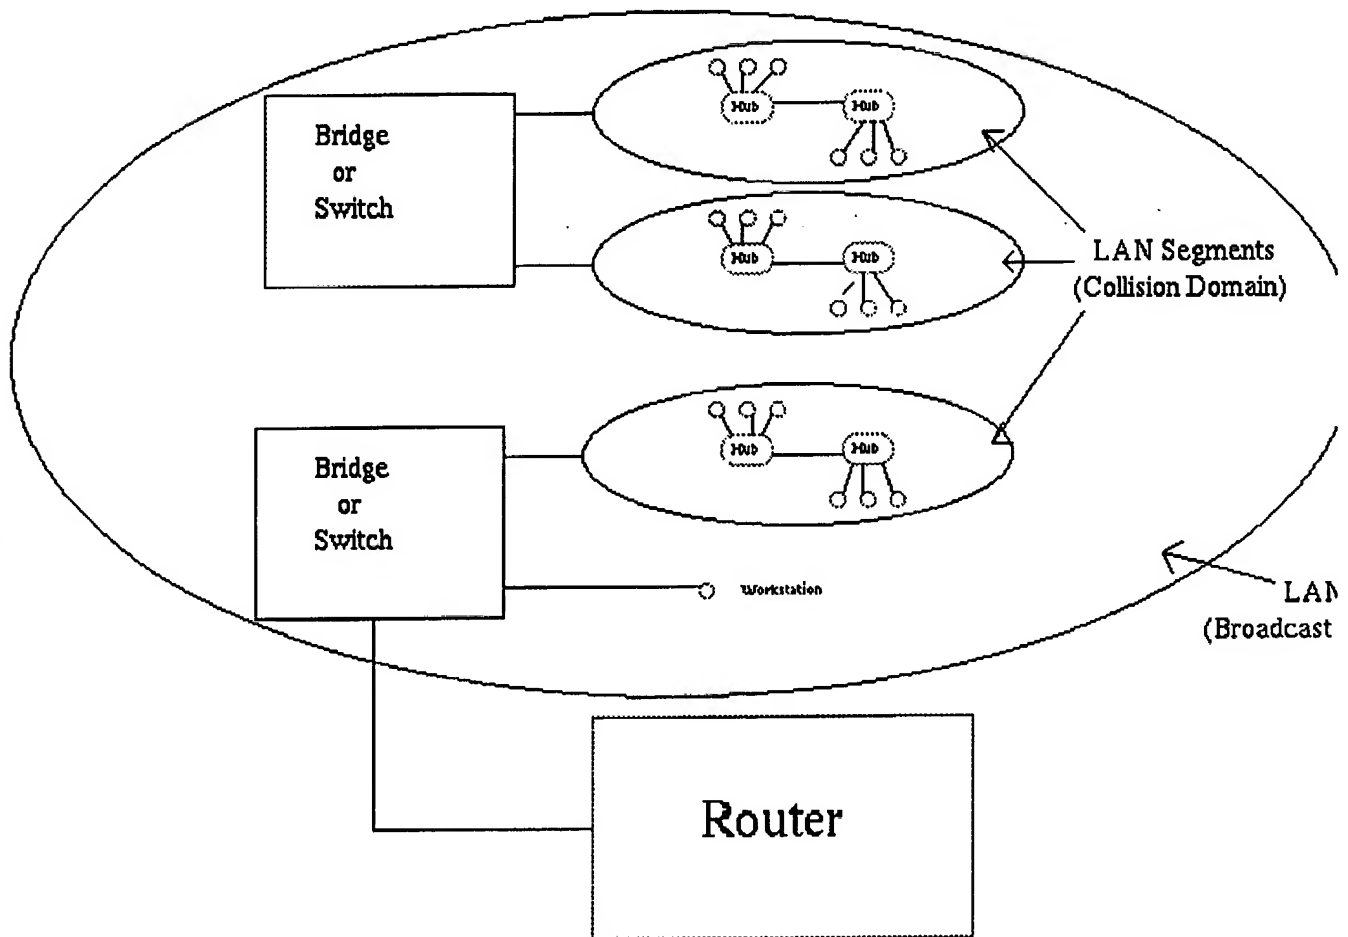


Figure 1: Physical view of a LAN.

VLAN's allow a network manager to logically segment a LAN into different broadcast domains (see *Figure 2*). Since this is a logical segmentation and not a physical one, workstations do not have to be physically located together. Users on different floors of the same building, or even in different buildings can now belong to the same LAN.

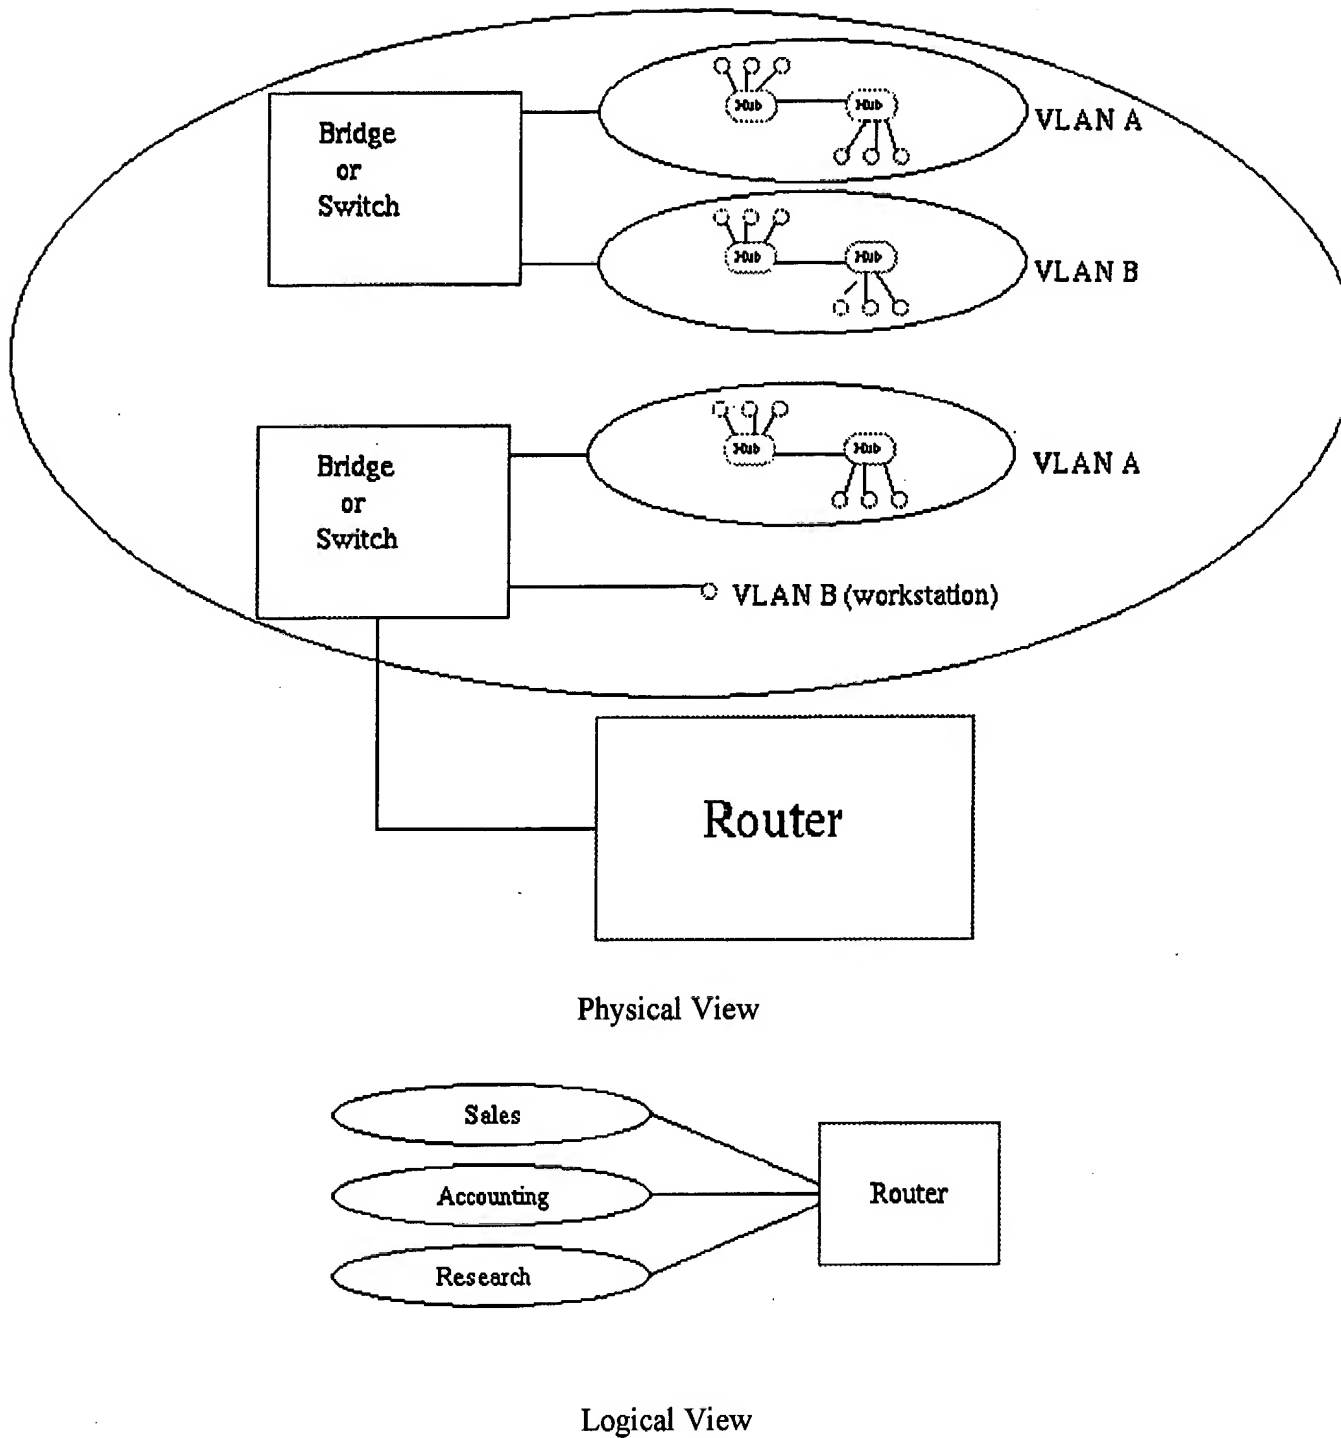


Figure 2: Physical and logical view of a VLAN.

VLAN's also allow broadcast domains to be defined without using routers. Bridging software is used instead to define which workstations are to be included in the broadcast domain. Routers would only have to be used to communicate between two VLAN's [Hein et al].

[Back to Table of Contents](#)

3.0 Why use VLAN's?

VLAN's offer a number of advantages over traditional LAN's. They are:

1) Performance

In networks where traffic consists of a high percentage of broadcasts and multicasts, VLAN's can reduce the need to send such traffic to unnecessary destinations. For example, in a broadcast domain consisting of 10 users, if the broadcast traffic is intended only for 5 of the users, then placing those 5 users on a separate VLAN can reduce traffic [[Passmore et al \(3Com report\)](#)].

Compared to switches, routers require more processing of incoming traffic. As the volume of traffic passing through the routers increases, so does the latency in the routers, which results in reduced performance. The use of VLAN's reduces the number of routers needed, since VLAN's create broadcast domains using switches instead of routers.

2) Formation of Virtual Workgroups

Nowadays, it is common to find cross-functional product development teams with members from different departments such as marketing, sales, accounting, and research. These workgroups are usually formed for a short period of time. During this period, communication between members of the workgroup will be high. To contain broadcasts and multicasts within the workgroup, a VLAN can be set up for them. With VLAN's it is easier to place members of a workgroup together. Without VLAN's, the only way this would be possible is to physically move all the members of the workgroup closer together.

However, virtual workgroups do not come without problems. Consider the situation where one user of the workgroup is on the fourth floor of a building, and the other workgroup members are on the second floor. Resources such as a printer would be located on the second floor, which would be inconvenient for the lone fourth floor user.

Another problem with setting up virtual workgroups is the implementation of centralized server farms, which are essentially collections of servers and major resources for operating a network at a central location. The advantages here are numerous, since it is more efficient and cost-effective to provide better security, uninterrupted power supply, consolidated backup, and a proper operating environment in a single area than if the major resources were scattered in a building. Centralized server farms can cause problems when setting up virtual workgroups if servers cannot be placed on more than one VLAN. In such a case, the server would be placed on a single VLAN and all other VLAN's trying to access the server would have to go through a router; this can reduce performance [Netreference Inc. article].

3) Simplified Administration

Seventy percent of network costs are a result of adds, moves, and changes of users in the network [[Buerger](#)]. Every time a user is moved in a LAN, recabling, new station addressing, and reconfiguration of hubs and routers becomes necessary. Some of these tasks can be simplified with the use of VLAN's. If a user is moved within a VLAN, reconfiguration of routers is unnecessary. In addition, depending on the type of VLAN, other administrative work can be reduced or eliminated [[Cisco white paper](#)]. However the full power of VLAN's will only really be felt when good management tools are created which can allow network managers to drag and drop

users into different VLAN's or to set up aliases.

Despite this saving, VLAN's add a layer of administrative complexity, since it now becomes necessary to manage virtual workgroups [[Passmore et al \(3Com report\)](#)].

4) Reduced Cost

VLAN's can be used to create broadcast domains which eliminate the need for expensive routers.

5) Security

Periodically, sensitive data may be broadcast on a network. In such cases, placing only those users who can have access to that data on a VLAN can reduce the chances of an outsider gaining access to the data. VLAN's can also be used to control broadcast domains, set up firewalls, restrict access, and inform the network manager of an intrusion [[Passmore et al \(3Com report\)](#)].

[Back to Table of Contents](#)

4.0 How VLAN's work

When a LAN bridge receives data from a workstation, it tags the data with a VLAN identifier indicating the VLAN from which the data came. This is called explicit tagging. It is also possible to determine to which VLAN the data received belongs using implicit tagging. In implicit tagging the data is not tagged, but the VLAN from which the data came is determined based on other information like the port on which the data arrived. Tagging can be based on the port from which it came, the source Media Access Control (MAC) field, the source network address, or some other field or combination of fields. VLAN's are classified based on the method used. To be able to do the tagging of data using any of the methods, the bridge would have to keep an updated database containing a mapping between VLAN's and whichever field is used for tagging. For example, if tagging is by port, the database should indicate which ports belong to which VLAN. This database is called a filtering database. Bridges would have to be able to maintain this database and also to make sure that all the bridges on the LAN have the same information in each of their databases. The bridge determines where the data is to go next based on normal LAN operations. Once the bridge determines where the data is to go, it now needs to determine whether the VLAN identifier should be added to the data and sent. If the data is to go to a device that knows about VLAN implementation (VLAN-aware), the VLAN identifier is added to the data. If it is to go to a device that has no knowledge of VLAN implementation (VLAN-unaware), the bridge sends the data without the VLAN identifier.

In order to understand how VLAN's work, we need to look at the types of VLAN's, the types of connections between devices on VLAN's, the filtering database which is used to send traffic to the correct VLAN, and tagging, a process used to identify the VLAN originating the data.

VLAN Standard: IEEE 802.1Q Draft Standard

There has been a recent move towards building a set of standards for VLAN products. The Institute of Electrical and Electronic Engineers (IEEE) is currently working on a draft standard 802.1Q for VLAN's. Up to this point, products have been proprietary, implying that anyone wanting to install VLAN's would have to purchase all products from the same vendor. Once the standards have been written and vendors

create products based on these standards, users will no longer be confined to purchasing products from a single vendor. The major vendors have supported these standards and are planning on releasing products based on them. It is anticipated that these standards will be ratified later this year.

[Back to Table of Contents](#)

4.1 Types of VLAN's

VLAN membership can be classified by port, MAC address, and protocol type.

1) Layer 1 VLAN: Membership by Port

Membership in a VLAN can be defined based on the ports that belong to the VLAN. For example, in a bridge with four ports, ports 1, 2, and 4 belong to VLAN 1 and port 3 belongs to VLAN 2 (see Figure3).

Port	VLAN
1	1
2	1
3	2
4	1

~~Port 1 - VLAN 1~~
VLAN 1 - Port 1, 2, 4

Figure3: Assignment of ports to different VLAN's.

The main disadvantage of this method is that it does not allow for user mobility. If a user moves to a different location away from the assigned bridge, the network manager must reconfigure the VLAN.

2) Layer 2 VLAN: Membership by MAC Address

Here, membership in a VLAN is based on the MAC address of the workstation. The switch tracks the MAC addresses which belong to each VLAN (see Figure4). Since MAC addresses form a part of the workstation's network interface card, when a workstation is moved, no reconfiguration is needed to allow the workstation to remain in the same VLAN. This is unlike Layer 1 VLAN's where membership tables must be reconfigured.

MAC Address	VLAN
1212354145121	1
2389234873743	2

3045834758445	2
5483573475843	1

Figure4: Assignment of MAC addresses to different VLAN's.

The main problem with this method is that VLAN membership must be assigned initially. In networks with thousands of users, this is no easy task. Also, in environments where notebook PC's are used, the MAC address is associated with the docking station and not with the notebook PC. Consequently, when a notebook PC is moved to a different docking station, its VLAN membership must be reconfigured.

3) Layer 2 VLAN: Membership by Protocol Type

VLAN membership for Layer 2 VLAN's can also be based on the protocol type field found in the Layer 2 header (see *Figure5*).

Protocol	VLAN
IP	1
IPX	2

Figure5: Assignment of protocols to different VLAN's.

4) Layer 3 VLAN: Membership by IP Subnet Address

Membership is based on the Layer 3 header. The network IP subnet address can be used to classify VLAN membership (see *Figure 6*).

IP Subnet	VLAN
23.2.24	1
26.21.35	2

Figure6: Assignment of IP subnet addresses to different VLAN's.

Although VLAN membership is based on Layer 3 information, this has nothing to do with network routing and should not be confused with router functions. In this method, IP addresses are used only as a mapping to determine membership in VLAN's. No other processing of IP addresses is done.

In Layer 3 VLAN's, users can move their workstations without reconfiguring their network

addresses. The only problem is that it generally takes longer to forward packets using Layer 3 information than using MAC addresses.

5) Higher Layer VLAN's

It is also possible to define VLAN membership based on applications or service, or any combination thereof. For example, file transfer protocol (FTP) applications can be executed on one VLAN and telnet applications on another VLAN.

The 802.1Q draft standard defines Layer 1 and Layer 2 VLAN's only. Protocol type based VLAN's and higher layer VLAN's have been allowed for, but are not defined in this standard. As a result, these VLAN's will remain proprietary.

[Back to Table of Contents](#)

4.2 Types of Connections

Devices on a VLAN can be connected in three ways based on whether the connected devices are VLAN-aware or VLAN-unaware. Recall that a VLAN-aware device is one which understands VLAN memberships (i.e. which users belong to a VLAN) and VLAN formats.

1) Trunk Link

All the devices connected to a trunk link, including workstations, must be VLAN-aware. All frames on a trunk link must have a special header attached. These special frames are called tagged frames (see *Figure 7*).

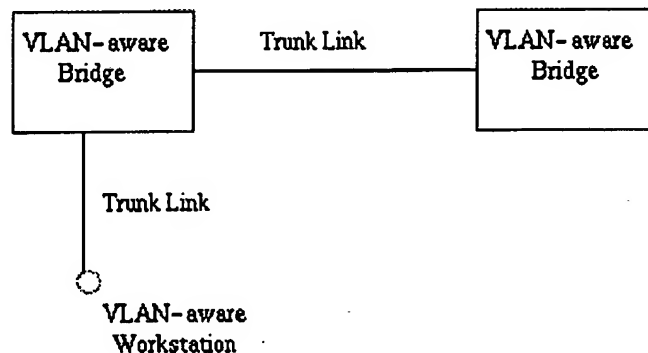


Figure 7: Trunk link between two VLAN-aware bridges.

2) Access Link

An access link connects a VLAN-unaware device to the port of a VLAN-aware bridge. All frames on access links must be implicitly tagged (untagged) (see *Figure 8*). The VLAN-unaware device can be a LAN segment with VLAN-unaware workstations or it can be a number of LAN segments containing VLAN-unaware devices (legacy LAN).



Figure 8: Access link between a VLAN-aware bridge and a VLAN-unaware device.

3) Hybrid Link

This is a combination of the previous two links. This is a link where both VLAN-aware and VLAN-unaware devices are attached (see Figure 9). A hybrid link can have both tagged and untagged frames, but *all* the frames for a specific VLAN must be either tagged or untagged.

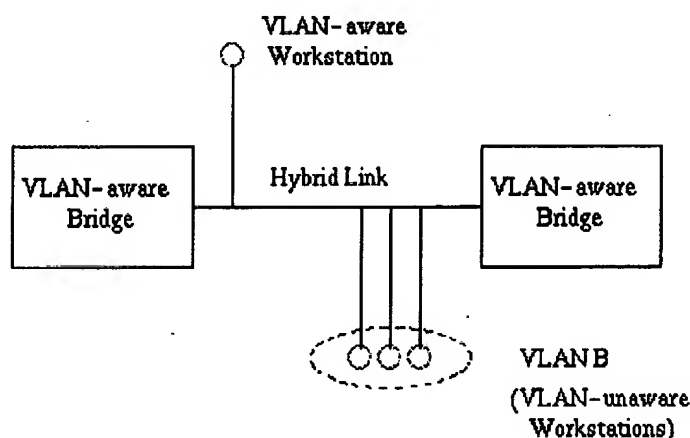


Figure 9: Hybrid link containing both VLAN-aware and VLAN-unaware devices.

It must also be noted that the network can have a combination of all three types of links.

[Back to Table of Contents](#)

4.3 Frame Processing

A bridge on receiving data determines to which VLAN the data belongs either by implicit or explicit tagging. In explicit tagging a tag header is added to the data. The bridge also keeps track of VLAN members in a filtering database which it uses to determine where the data is to be sent. Following is an explanation of the contents of the filtering database and the format and purpose of the tag header [802.1Q].

1) Filtering Database

Membership information for a VLAN is stored in a filtering database. The filtering database consists of the following types of entries:

i) Static Entries

Static information is added, modified, and deleted by management only. Entries are not automatically removed after some time (ageing), but must be explicitly removed by management. There are two types of static entries:

- a) Static Filtering Entries: which specify for every port whether frames to be sent to a specific MAC address or group address and on a specific VLAN should be forwarded or discarded, or should follow the dynamic entry, and
- b) Static Registration Entries: which specify whether frames to be sent to a specific VLAN are to be tagged or untagged and which ports are registered for that VLAN.

ii) Dynamic Entries

Dynamic entries are learned by the bridge and cannot be created or updated by management. The learning process observes the port from which a frame, with a given source address and VLAN ID (VID), is received, and updates the filtering database. The entry is updated only if all the following three conditions are satisfied:

- a) this port allows learning,
- b) the source address is a workstation address and not a group address, and
- c) there is space available in the database.

Entries are removed from the database by the ageing out process where, after a certain amount of time specified by management (10 sec --- 1000000 sec), entries allow automatic reconfiguration of the filtering database if the topology of the network changes. There are three types of dynamic entries:

- a) Dynamic Filtering Entries: which specify whether frames to be sent to a specific MAC address and on a certain VLAN should be forwarded or discarded.
- * b) Group Registration Entries: which indicate for each port whether frames to be sent to a group MAC address and on a certain VLAN should be filtered or discarded. These entries are added and deleted using Group Multicast Registration Protocol (GMRP). This allows multicasts to be sent on a single VLAN without affecting other VLAN's.
- c) Dynamic Registration Entries: which specify which ports are registered for a specific VLAN. Entries are added and deleted using GARP VLAN Registration Protocol (GVRP), where GARP is the Generic Attribute Registration Protocol.

GVRP is used not only to update dynamic registration entries, but also to communicate the information to other VLAN-aware bridges.

In order for VLAN's to forward information to the correct destination, all the bridges in the VLAN should contain the same information in their respective filtering databases. GVRP allows both VLAN-aware workstations and bridges to issue and revoke VLAN memberships. VLAN-aware bridges register and propagate VLAN membership to all ports that are a part of the active topology of the VLAN. The active topology of a network is determined when the bridges are

turned on or when a change in the state of the current topology is perceived.

The active topology is determined using a spanning tree algorithm which prevents the formation of loops in the network by disabling ports. Once an active topology for the network (which may contain several VLAN's) is obtained, the bridges determine an active topology for each VLAN. This may result in a different topology for each VLAN or a common one for several VLAN's. In either case, the VLAN topology will be a subset of the active topology of the network (see *Figure 10*).

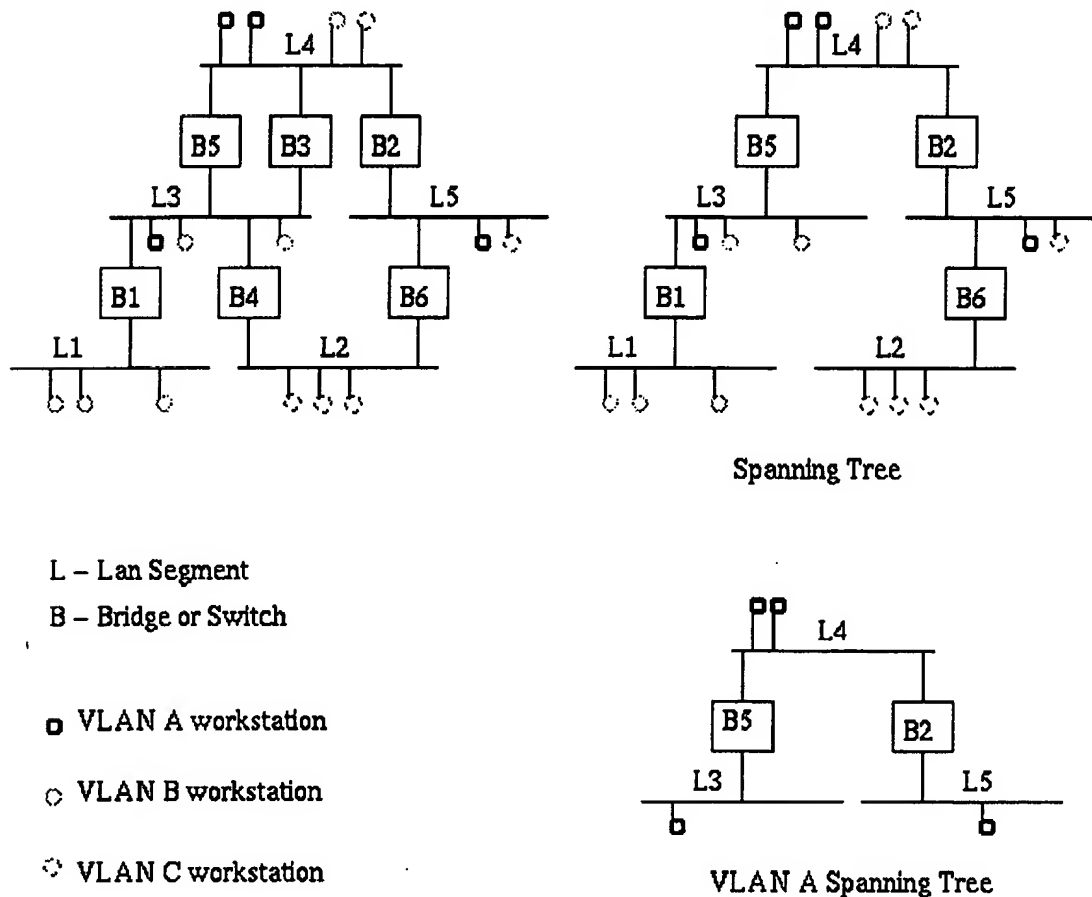


Figure 10: Active topology of network and VLAN A using spanning tree algorithm.

2) Tagging

When frames are sent across the network, there needs to be a way of indicating to which VLAN the frame belongs, so that the bridge will forward the frames only to those ports that belong to that VLAN, instead of to all output ports as would normally have been done. This information is added to the frame in the form of a tag header. In addition, the tag header:

- i) allows user priority information to be specified,
- ii) allows source routing control information to be specified, and

iii) indicates the format of MAC addresses.

Frames in which a tag header has been added are called tagged frames. Tagged frames convey the VLAN information across the network.

The tagged frames that are sent across hybrid and trunk links contain a tag header. There are two formats of the tag header:

i) Ethernet Frame Tag Header: The ethernet frame tag header (see *Figure 11*) consists of a tag protocol identifier (TPID) and tag control information (TCI).




Figure 11: Ethernet frame tag header.

ii) Token Ring and Fiber Distributed Data Interface (FDDI) tag header: The tag headers for both token ring and FDDI networks consist of a SNAP-encoded TPID and TCI.




Figure 12: Token ring and FDDI tag header.

TPID is the tag protocol identifier which indicates that a tag header is following and TCI (see *Figure 13*) contains the user priority, canonical format indicator (CFI), and the VLAN ID.



Figure 13: Tag control information (TCI).

User priority is a 3 bit field which allows priority information to be encoded in the frame. Eight levels of priority are allowed, where zero is the lowest priority and seven is the highest priority. How this field is used is described in the supplement 802.1p.

The CFI bit is used to indicate that all MAC addresses present in the MAC data field are in canonical format. This field is interpreted differently depending on whether it is an ethernet-encoded tag header or a SNAP-encoded tag header. In SNAP-encoded TPID the field indicates the presence or absence of the canonical format of addresses. In ethernet-encoded TPID, it indicates the presence of the Source-Routing Information (RIF) field after the length field. The RIF field indicates routing on ethernet frames.

The VID field is used to uniquely identify the VLAN to which the frame belongs. There can be a maximum of $(2^{12} - 1)$ VLAN's. Zero is used to indicate no VLAN ID, but that user priority information is present. This allows priority to be encoded in non-priority LAN's.

[Back to Table of Contents](#)

5.0 Summary

As we have seen there are significant advances in the field of networks in the form of VLAN's which allow the formation of virtual workgroups, better security, improved performance, simplified administration, and reduced costs. VLAN's are formed by the logical segmentation of a network and can be classified into Layer 1, 2, 3 and higher layers. Only Layer 1 and 2 are specified in the draft standard 802.1Q. Tagging and the filtering database allow a bridge to determine the source and destination VLAN for received data. VLAN's if implemented effectively, show considerable promise in future networking solutions.

[Back to Table of Contents](#)

6.0 References

- 1) David Passmore, John Freeman, "The Virtual LAN Technology Report," March 7, 1997,
<http://www.3com.com/nsc/200374.html>
A very good overview of VLAN's, their strengths, weaknesses, and implementation problems.
- 2) IEEE, "Draft Standard for Virtual Bridge Local Area Networks," P802.1Q/D1, May 16, 1997,
This is the draft standard for VLAN's which covers implementation issues of Layer 1 and 2 VLAN's.
- 3) Mathias Hein, David Griffiths, Orna Berry, "Switching Technology in the Local Network: From LAN to Switched LAN to Virtual LAN," February 1997,
Textbook explanation of what VLAN's are and their types.
- 7) Susan Biagi, "Virtual LANs," Network VAR v4 n1 p. 10-12, January 1996,
An Overview of VLAN's, advantages, and disadvantages.
- 8) David J. Buerger, "Virtual LAN cost savings will stay virtual until networking's next era," Network World, March 1995,
A short summary on VLAN's.
- 9) IEEE, "Traffic Class Expediting and Dynamic Multicast Filtering," 802.1p/D6, April 1997,
This is the standard for implementing priority and dynamic multicasts. Implementation of priority in VLAN's is based on this standard.

[Back to Table of Contents](#)

7.0 Abbreviations

- CFI - Canonical Format Indicator
- FDDI - Fiber Distributed Data Interface
- FTP - File Transfer Protocol
- GARP - Generic Attribute Registration Protocol
- GMRP - Group Multicast Registration Protocol
- GVRP - GARP VLAN Registration Protocol
- IEEE - Institute of Electrical and Electronic Engineers
- LAN - Local Area Network
- MAC - Media Access Control
- RIF - Source-Routing Information
- TCI - Tag Control Information
- TPID - Tag Protocol Identifier
- VID - VLAN ID
- VLAN - Virtual Local Area Network

Back to [Table of Contents](#)

Last Modified: August 14, 1997

Ref 2

[<<< Back](#) | [Print](#)

Ternary DRAM CAM: Now and the Future
 Devashish Paul, MOSAID Semiconductor -- 5/1/2000
 Electronic News


In the past year the data networking industry has shown a growing interest in **Ternary Content Addressable Memories (CAMs)**, especially those that use DRAM technology due to the basic efficiency of the DRAM CAM cell.

With the explosion in Internet usage and corresponding bandwidth growth, every step in packet routing or cell switching must be optimized, including table lookup and packet classification. Ternary CAMs can store 1,0, and X (don't care values) and are ideally suited to accelerate the table lookup function performed in today's networking applications. Unlike conventional memories, which return the content of a given memory address, a CAM returns the address of a given cell containing the content of a word that is being searched, an ideal capability for high-speed lookups of large internetworking tables.

Advertisement

*Devashish Paul is
 a product
 marketing
 manager for
 MOSAID
 Semiconductor.*

Unlike software-based table-lookup implementations, CAMs can perform table lookups at wire speed, conducting as many as 66 million searches per second, an increase of several orders of magnitude in performance. Ternary CAMs are typically targeted at data communications applications such as carrier class edge and enterprise routers, core Internet Gigabit and Terabit routers, SONET systems, and asynchronous transfer mode switching equipment. They are critical for a number of Internetworking applications such as Classless Inter Domain Routing, flow analysis, advanced VLAN support, and policy-based networking (QOS/TOS/COS), where a search only needs to be performed on part of a word.



A Ternary DRAM CAM allows for a denser CAM with a much smaller cell size than SRAM-based solutions. DRAM CAM cells can be implemented in just six transistors, representing a 2.5 times density advantage over competing SRAM in the same process technology generation.

Fixed overhead logic allows DRAM CAMs to scale to increasingly wide word implementations with only minimal die size penalties. The ability to create cost-effective, extremely wide CAMs is critical for emerging applications such as flow analysis, RSVP, and IPv6.

As semiconductor manufacturing processes evolve over the next few years, it will be possible to design and implement the deeper and wider CAMs that the networking OEMs require. Currently, the largest

available densities are 2Mbit devices with word widths of 288 bits and a corresponding table depth of 8Kbit entries. To achieve larger tables, networking OEMs must cascade a number of CAMs. The preferred solution is to use one large chip with a maximum amount of memory. Some Ternary CAM vendors are currently announcing 4Mbit and even 8Mbit or 9Mbit devices to meet the networking market's need for bigger and wider CAMs.

Designing and manufacturing large monolithic CAMs is limited by the maximum practical die size in a given process technology. As process technologies mature, designers can pack more memory and logic into the same die size. This gives the DRAM-based CAM a large advantage over SRAM CAMs since its cell size is 2.5 times smaller than a corresponding SRAM cell. In today's technology it is already possible to implement Ternary DRAM CAMs as large as 9Mbits. With 0.15-micron and 0.13-micron merged DRAM-logic processes becoming available, CAMs ranging from 18Mbit to 36Mbit will become feasible.

Today's DRAM-based Ternary CAMs provide the baseline technology for tomorrow's large and wide CAMs required by the most advanced networking applications. DRAM Ternary CAMs use a 6T cell that uses less area and scales better than the 16T SRAM CAM cell for the implementation of the large and wide Ternary CAMs of the future. DRAM-based CAMs are poised to become the mainstream CAM technology of the future. MOSAID Semiconductor has already launched the industry's first DRAM-based CAM with a density of 2.3Mbit and will launch even larger and wider CAMs in the future. Other vendors are also entering the market, basing future CAMs on DRAM technology that is ideally suited to the networking industry's need for large, deep, and wide CAMs.

Devashish Paul is a product marketing manager for MOSAID Semiconductor.

[<<< Back](#) | [Print](#)

© 2005, Reed Business Information, a division of Reed Elsevier Inc. All Rights Reserved.

RANDOLPH

P.O. BOX 1450

ALEXANDRIA, VA 22313-1450

IF UNDELIVERABLE RETURN IN TEN DAYS

OFFICIAL BUSINESS

AN EQUAL OPPORTUNITY EMPLOYER



02 1A
0004205065
MAILED FROM ZIP CODE
\$0.
JUN
UNITED STATES POSTAGE
PENAL
PRIVACY
FIRST CLASS
MONEY